

10/670,561 STIC/EIC search

Set	Items	Description
S1	4509	DATABASE? ? OR DBMS OR RDBMS OR OODB OR DATA()BASE? ?
S2	444	AGGREGATION? ? OR AGGREGATE? ? OR AGGREGATING
S3	8695	GROUP? ? OR SET? ? OR SUBSET? ? OR COLLECTION? ? OR ARRAY?
	?	
S4	46	(S2 OR S3) (5N) (UPDATE? ? OR UPDATING OR UP() (DATE? ? OR - DATING) OR RENEW? ? OR RENEWED OR RENEWING OR REVISION? ? OR - REVI?E? ?)
S5	170	PARTITION?? OR PARTITIONING
S6	5629	TRANSACTION? ? OR ACTIVITY OR ACTIVITIES OR OPERATION? ?
S7	346	S1 (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL - OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S8	451	(S2 OR S3) (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR - SEVERAL OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S9	14	S5 (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL - OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S10	3081	INSTANCE? ?
S11	558	(S10 OR S6) (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S12	0	S4 AND S7 AND S8 AND S9 AND S11
S13	0	S7 AND S8 AND S9 AND S11 AND (UPDATE? ? OR UPDATING OR UP(-) (DATE? ? OR DATING) OR RENEW? ? OR RENEWED OR RENEWING OR RE- VISION? ? OR REVI?E? ?)
S14	0	S4 AND S1 AND S5 AND (S6 OR S10)

? show files

File 256:TecInfoSource 82-2006/Dec

(c) 2006 Info.Sources Inc

Set	Items	Description
S1	330075	DATABASE? ? OR DBMS OR RDBMS OR OODB OR DATA()BASE? ?
S2	177549	AGGREGATION? ? OR AGGREGATE? ? OR AGGREGATING
S3	3760245	GROUP? ? OR SET? ? OR SUBSET? ? OR COLLECTION? ? OR ARRAY?
		?
S4	19571	(S2 OR S3) (5N) (UPDATE? ? OR UPDATING OR UP() (DATE? ? OR - DATING) OR RENEW? ? OR RENEWED OR RENEWING OR REVISION? ? OR - REVI?E? ?)
S5	241085	PARTITION?? OR PARTITIONING
S6	3040167	TRANSACTION? ? OR ACTIVITY OR ACTIVITIES OR OPERATION? ?
S7	42130	S1 (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL - OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S8	564402	(S2 OR S3) (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR - SEVERAL OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S9	25944	S5 (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL - OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S10	475673	INSTANCE? ?
S11	316763	(S10 OR S6) (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S12	0	S4 (30N) S7 (30N) S8 (30N) S9 (30N) S11
S13	1	S4 (50N) S7 (50N) S8 (50N) S9 (50N) S11
S14	6	(UPDATE? ? OR UPDATING OR UP() (DATE? ? OR DATING) OR RENEW? ? OR RENEWED OR RENEWING OR REVISION? ? OR REVI?E? ?) (50N) - S7 (50N) S8 (50N) S9 (50N) S11
S15	5	S14 NOT S13
S16	5	IDPAT (sorted in duplicate/non-duplicate order)
S17	5	IDPAT (primary/non-duplicate records only)
S18	27	S4 (30N) S1 (30N) S5 (30N) (S6 OR S10)
S19	26	S18 NOT (S13 OR S17)
S20	26	IDPAT (sorted in duplicate/non-duplicate order)
S21	26	IDPAT (primary/non-duplicate records only)
S22	24	S4 (30N) S1 (30N) S5 (30N) S6
S23	23	S22 NOT (S13 OR S17)
S24	22	S23 AND IC=G06F
S25	22	IDPAT (sorted in duplicate/non-duplicate order)
S26	22	IDPAT (primary/non-duplicate records only)
S27	1655	S4 (30N) (S7 OR S8 OR S9 OR S11)

File 348:EUROPEAN PATENTS 1978-2006/ 200636

(c) 2006 European Patent Office

File 349:PCT FULLTEXT 1979-2006/UB=20060831UT=20060824

(c) 2006 WIPO/Thomson

File 350:Derwent WPIX 1963-2006/UD=200657

(c) 2006 The Thomson Corporation

Set	Items	Description
S1	873155	DATABASE? ? OR DBMS OR RDBMS OR OODB OR DATA()BASE? ?
S2	486304	AGGREGATION? ? OR AGGREGATE? ? OR AGGREGATING
S3	8962816	GROUP? ? OR SET? ? OR SUBSET? ? OR COLLECTION? ? OR ARRAY?
S4	15592	(S2 OR S3) (5N) (UPDATE? ? OR UPDATING OR UP() (DATE? ? OR - DATING) OR RENEW? ? OR RENEWED OR RENEWING OR REVISION? ? OR - REVI?E? ?)
S5	319052	PARTITION?? OR PARTITIONING
S6	7371553	TRANSACTION? ? OR ACTIVITY OR ACTIVITIES OR OPERATION? ?
S7	29471	S1 (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL - OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S8	295854	(S2 OR S3) (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR - SEVERAL OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S9	11003	S5 (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL - OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S10	201929	INSTANCE? ?
S11	205733	(S10 OR S6) (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S12	0	S4 AND S7 AND S8 AND S9 AND S11
S13	0	S7 AND S8 AND S9 AND S11 AND (UPDATE? ? OR UPDATING OR UP(-) (DATE? ? OR DATING) OR RENEW? ? OR RENEWED OR RENEWING OR RE- VISION? ? OR REVI?E? ?)
S14	16	S4 AND S1 AND S5 AND (S6 OR S10)
S15	16	S14 NOT PY>2003
S16	14	RD (unique items)
File	8: Ei Compendex(R) 1970-2006/Sep W1	(c) 2006 Elsevier Eng. Info. Inc.
File	35: Dissertation Abs Online 1861-2006/Aug	(c) 2006 ProQuest Info&Learning
File	65: Inside Conferences 1993-2006/Sep 11	(c) 2006 BLDSC all rts. reserv.
File	2: INSPEC 1898-2006/Sep W1	(c) 2006 Institution of Electrical Engineers
File	94: JICST-EPlus 1985-2006/Jun W1	(c) 2006 Japan Science and Tech Corp(JST)
File	111: TGG Natl. Newspaper Index(SM) 1979-2006/Aug 28	(c) 2006 The Gale Group
File	6: NTIS 1964-2006/Sep W1	(c) 2006 NTIS, Intl Cpyrght All Rights Res
File	144: Pascal 1973-2006/Aug W3	(c) 2006 INIST/CNRS
File	434: SciSearch(R) Cited Ref Sci 1974-1989/Dec	(c) 2006 The Thomson Corp
File	34: SciSearch(R) Cited Ref Sci 1990-2006/Sep W1	(c) 2006 The Thomson Corp
File	62: SPIN(R) 1975-2006/Aug W4	(c) 2006 American Institute of Physics
File	99: Wilson Appl. Sci & Tech Abs 1983-2006/Jul	(c) 2006 The HW Wilson Co.
File	95: TEME-Technology & Management 1989-2006/Sep W2	(c) 2006 FIZ TECHNIK
File	56: Computer and Information Systems Abstracts 1966-2006/Aug	(c) 2006 CSA.
File	57: Electronics & Communications Abstracts 1966-2006/Aug	(c) 2006 CSA.
File	60: ANTE: Abstracts in New Tech & Engineer 1966-2006/Aug	(c) 2006 CSA.
File	266: FEDRIP 2006/Aug	Comp & dist by NTIS, Intl Copyright All Rights Res
File	583: Gale Group Globalbase(TM) 1986-2002/Dec 13	

(c) 2002 The Gale Group
File 438:Library Lit. & Info. Science 1984-2006/Aug
(c) 2006 The HW Wilson Co

Set	Items	Description
S1	53776	DATABASE? ? OR DBMS OR RDBMS OR OODB OR DATA()BASE? ?
S2	28019	AGGREGATION? ? OR AGGREGATE? ? OR AGGREGATING
S3	1348114	GROUP? ? OR SET? ? OR SUBSET? ? OR COLLECTION? ? OR ARRAY?
		?
S4	2749	(S2 OR S3) (5N) (UPDATE? ? OR UPDATING OR UP() (DATE? ? OR - DATING) OR RENEW? ? OR RENEWED OR RENEWING OR REVISION? ? OR - REVI?E? ?)
S5	82184	PARTITION?? OR PARTITIONING
S6	1112360	TRANSACTION? ? OR ACTIVITY OR ACTIVITIES OR OPERATION? ?
S7	2101	S1 (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL - OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S8	79422	(S2 OR S3) (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR - SEVERAL OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S9	5746	S5 (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL - OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S10	46013	INSTANCE? ?
S11	24816	(S10 OR S6) (3N) (PLURAL? OR MORE OR LEAST()ONE OR MANY OR SEVERAL OR MULTIPLE? ? OR MULTIPLICITY OR MULTI OR VARIOUS OR NUMEROUS OR NUMBER()OF OR MYRIAD OR THOUSANDS OR MILLIONS)
S12	0	S4 AND S7 AND S8 AND S9 AND S11
S13	0	S7 AND S8 AND S9 AND S11 AND (UPDATE? ? OR UPDATING OR UP(-) (DATE? ? OR DATING) OR RENEW? ? OR RENEWED OR RENEWING OR RE- VISION? ? OR REVI?E? ?)
S14	0	S4 AND S1 AND S5 AND (S6 OR S10)

File 347:JAPIO Dec 1976-2005/Dec(Updated 060404)
(c) 2006 JPO & JAPIO

☐ Search Results

[BROWSE](#)

[SEARCH](#)

[IEEE XPLORE GUIDE](#)

Results for "(((database<in>metadata) <and> (partition<in>metadata))<and> (olap<..."

Your search matched 6 of 1408155 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

☒ e-mail

» Search Options

[View Session History](#)

[New Search](#)

Modify Search

(((database<in>metadata) <and> (partition<in>metadata))<and> (olap<in>metad

[Search](#)

☐ Check to search only within this results set

Display Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

[view selected items](#)

[Select All](#) [Deselect All](#)

- ☐ 1. **OLAP query processing for partitioned data warehouses**
Bellatreche, L.; Karlapalem, K.; Mohania, M.;
[Database Applications in Non-Traditional Environments, 1999. \(DANTE '99\) Pr International Symposium on](#)
1999 Page(s):35 - 42
Digital Object Identifier 10.1109/DANTE.1999.844939
[AbstractPlus](#) | Full Text: [PDF](#)(312 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 2. **Exploitation of pre-sortedness for sorting in query processing: the Temp for UB-trees**
Zirkel, M.; Markl, V.; Bayer, R.;
[Database Engineering & Applications, 2001 International Symposium on. 16-18 July 2001 Page\(s\):155 - 166](#)
Digital Object Identifier 10.1109/IDEAS.2001.938082
[AbstractPlus](#) | Full Text: [PDF](#)(1044 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 3. **What can partitioning do for your data warehouses and data marts?**
Bellatreche, L.; Karlapalem, K.; Mohania, M.; Schneider, M.;
[Database Engineering and Applications Symposium, 2000 International 18-20 Sept. 2000 Page\(s\):437 - 445](#)
Digital Object Identifier 10.1109/IDEAS.2000.880634
[AbstractPlus](#) | Full Text: [PDF](#)(720 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 4. **Management of multidimensional aggregates for efficient online analytics**
Albrecht, J.; Bauer, A.; Deyerling, O.; Gunzel, H.; Hummer, Lehner, W.; Schles
[Database Engineering and Applications, 1999. IDEAS '99. International Sympo Proceedings](#)
2-4 Aug. 1999 Page(s):156 - 164
Digital Object Identifier 10.1109/IDEAS.1999.787264
[AbstractPlus](#) | Full Text: [PDF](#)(144 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 5. **A case for parallelism in data warehousing and OLAP**
Datta, A.; Bongki Moon; Thomas, H.;
[Database and Expert Systems Applications, 1998. Proceedings, Ninth Internat](#)

☐ Search Results

[BROWSE](#)

[SEARCH](#)

[IEEE XPLORE GUIDE](#)

Results for "(((aggregates<in>metadata) <and> (sort<in>metadata))<and> (database<in>me

 e-mail

Your search matched 4 of 1408155 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

[View Session History](#)

[New Search](#)

Modify Search

(((aggregates<in>metadata) <and> (sort<in>metadata))<and> (database<in>me

[Search](#)

☐ Check to search only within this results set

Display Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

[view selected items](#)

[Select All](#) [Deselect All](#)

- ☐ 1. **Multi-cube computation**
Xu Yu, J.; Hongjun Lu;
[Database Systems for Advanced Applications, 2001. Proceedings. Seventh International Conference on](#)
18-21 April 2001 Page(s):126 - 133
Digital Object Identifier 10.1109/DASFAA.2001.916373
[AbstractPlus](#) | Full Text: [PDF](#)(636 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 2. **Computing temporal aggregates**
Kline, N.; Snodgrass, R.T.;
[Data Engineering, 1995. Proceedings of the Eleventh International Conference](#)
6-10 March 1995 Page(s):222 - 231
Digital Object Identifier 10.1109/ICDE.1995.380389
[AbstractPlus](#) | Full Text: [PDF](#)(916 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 3. **Relational division: four algorithms and their performance**
Graefe, G.;
[Data Engineering, 1989. Proceedings. Fifth International Conference on](#)
6-10 Feb. 1989 Page(s):94 - 101
Digital Object Identifier 10.1109/ICDE.1989.47204
[AbstractPlus](#) | Full Text: [PDF](#)(724 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 4. **Informix parallel data query (PDQ)**
Clay, D.;
[Parallel and Distributed Information Systems, 1993.. Proceedings of the Second Conference on](#)
20-22 Jan. 1993 Page(s):71 - 72
Digital Object Identifier 10.1109/PDIS.1993.253069
[AbstractPlus](#) | Full Text: [PDF](#)(148 KB) IEEE CNF
[Rights and Permissions](#)

[on](#)

26-28 Aug. 1998 Page(s):226 - 231

Digital Object Identifier 10.1109/DEXA.1998.707407

[AbstractPlus](#) | [Full Text: PDF\(288 KB\)](#) [IEEE CNF](#)
[Rights and Permissions](#)



6. Parallel Star Join+DataIndexes: efficient query processing in data wareh

Datta, A.; VanderMeer, D.; Ramamritham, K.;

[Knowledge and Data Engineering, IEEE Transactions on](#)

Volume 14, Issue 6, Nov.-Dec. 2002 Page(s):1299 - 1316

Digital Object Identifier 10.1109/TKDE.2002.1047769

[AbstractPlus](#) | [References](#) | [Full Text: PDF\(1973 KB\)](#) [IEEE JNL](#)
[Rights and Permissions](#)

Indexed by



[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE –

☐ Search Results

[BROWSE](#)

[SEARCH](#)

[IEEE XPLORE GUIDE](#)

Results for "((aggregation<in>metadata) <and> (database<in>metadata))<and> (partit..."

Your search matched **17** of **1408155** documents.

A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

☒ e-mail

» Search Options

[View Session History](#)

[New Search](#)

» Key

IEEE JNL IEEE Journal or Magazine
IEE JNL IEE Journal or Magazine
IEEE CNF IEEE Conference Proceeding
IEE CNF IEE Conference Proceeding
IEEE STD IEEE Standard

Modify Search

((aggregation<in>metadata) <and> (database<in>metadata))<and> (partition<in>

☐ Check to search only within this results set

Display Format: ☒ Citation ☐ Citation & Abstract

[Select All](#) [Deselect All](#)

- ☐ 1. **Mathematical model of composite objects and its application for organizing databases**
Ketabchi, M.A.; Berzins, V.;
[Software Engineering, IEEE Transactions on](#)
Volume 14, Issue 1, Jan. 1988 Page(s):71 - 84
Digital Object Identifier 10.1109/32.4624
[AbstractPlus](#) | Full Text: [PDF](#)(1300 KB) IEEE JNL
[Rights and Permissions](#)
- ☐ 2. **A generic algorithmic framework for aggregation of spatio-temporal data**
Seung-Hyun Jeong; Fernandes, A.A.A.; Paton, N.W.; Griffiths, T.;
[Scientific and Statistical Database Management, 2004. Proceedings. 16th International Conference on](#)
21-23 June 2004 Page(s):245 - 254
Digital Object Identifier 10.1109/SSDM.2004.1311216
[AbstractPlus](#) | Full Text: [PDF](#)(392 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 3. **Exploitation of pre-sortedness for sorting in query processing: the Temp for UB-trees**
Zirkel, M.; Markl, V.; Bayer, R.;
[Database Engineering & Applications, 2001 International Symposium on](#)
16-18 July 2001 Page(s):155 - 166
Digital Object Identifier 10.1109/IDEAS.2001.938082
[AbstractPlus](#) | Full Text: [PDF](#)(1044 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 4. **Parallel algorithms for computing temporal aggregates**
Gendrano, J.A.G.; Huang, B.C.; Rodrigue, J.M.; Bongki Moon; Snodgrass, R.T
[Data Engineering, 1999. Proceedings., 15th International Conference on](#)
23-26 March 1999 Page(s):418 - 427
Digital Object Identifier 10.1109/ICDE.1999.754958
[AbstractPlus](#) | Full Text: [PDF](#)(164 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 5. **Reuse of high-level information requests: leveraging the investment**
Washburn, G.A.; Delcambre, L.M.L.; Whiting, M.A.;
[Scientific and Statistical Database Systems, 1996. Proceedings., Eighth International Conference on](#)

Conference on

18-20 June 1996 Page(s):186 - 195

Digital Object Identifier 10.1109/SSDM.1996.506061

[AbstractPlus](#) | Full Text: [PDF](#)(884 KB) IEEE CNF

[Rights and Permissions](#)

- ☐ **6. Dynamic Load Balancing for the Distributed Mining of Molecular Structures**
Di Fatta, G.; Berthold, M.R.;
[Parallel and Distributed Systems, IEEE Transactions on](#)
Volume 17, Issue 8, Aug. 2006 Page(s):773 - 785
Digital Object Identifier 10.1109/TPDS.2006.101
[AbstractPlus](#) | Full Text: [PDF](#)(1720 KB) IEEE JNL
[Rights and Permissions](#)

- ☐ **7. Materialized Views in the Presence of Reporting Functions**
Habich, D.; Lehner, W.; Just, M.;
[Scientific and Statistical Database Management, 2006. 18th International Conference on](#)
03-05 July 2006 Page(s):159 - 168
Digital Object Identifier 10.1109/SSDBM.2006.32
[AbstractPlus](#) | Full Text: [PDF](#)(592 KB) IEEE CNF
[Rights and Permissions](#)

- ☐ **8. Evaluation of various aggregation operators applied to a content image retrieval**
Ionescu, M.; Ralescu, A.;
[Fuzzy Information Processing Society, 2005. NAFIPS 2005. Annual Meeting of the IEEE](#)
26-28 June 2005 Page(s):670 - 675
Digital Object Identifier 10.1109/NAFIPS.2005.1548618
[AbstractPlus](#) | Full Text: [PDF](#)(3784 KB) IEEE CNF
[Rights and Permissions](#)

- ☐ **9. MM-Cubing: computing Iceberg cubes by factorizing the lattice space**
Zheng Shao; Jiawei Han; Dong Xin;
[Scientific and Statistical Database Management, 2004. Proceedings. 16th International Conference on](#)
21-23 June 2004 Page(s):213 - 222
Digital Object Identifier 10.1109/SSDM.2004.1311213
[AbstractPlus](#) | Full Text: [PDF](#)(347 KB) IEEE CNF
[Rights and Permissions](#)

- ☐ **10. Securing OLAP data cubes against privacy breaches**
Lingyu Wang; Jajodia, S.; Wijesekera, D.;
[Security and Privacy, 2004. Proceedings. 2004 IEEE Symposium on](#)
9-12 May 2004 Page(s):161 - 175
Digital Object Identifier 10.1109/SECPRI.2004.1301322
[AbstractPlus](#) | Full Text: [PDF](#)(1520 KB) IEEE CNF
[Rights and Permissions](#)

- ☐ **11. On some fuzzy extensions of association rules**
Bosc, P.; Pivert, O.;
[IFSA World Congress and 20th NAFIPS International Conference, 2001. Joint](#)
Volume 2, 25-28 July 2001 Page(s):1104 - 1109 vol.2
Digital Object Identifier 10.1109/NAFIPS.2001.944759
[AbstractPlus](#) | Full Text: [PDF](#)(428 KB) IEEE CNF
[Rights and Permissions](#)

- ☐ **12. What can partitioning do for your data warehouses and data marts?**
Bellatreche, L.; Karlapalem, K.; Mohania, M.; Schneider, M.;
[Database Engineering and Applications Symposium, 2000 International](#)

18-20 Sept. 2000 Page(s):437 - 445
Digital Object Identifier 10.1109/IDEAS.2000.880634
[AbstractPlus](#) | [Full Text: PDF\(720 KB\)](#) IEEE CNF
[Rights and Permissions](#)

- ☐ **13. A cost function for uniformly partitioned UB-trees**
Markl, V.; Bayer, R.;
[Database Engineering and Applications Symposium, 2000 International](#)
18-20 Sept. 2000 Page(s):410 - 416
Digital Object Identifier 10.1109/IDEAS.2000.880626
[AbstractPlus](#) | [Full Text: PDF\(672 KB\)](#) IEEE CNF
[Rights and Permissions](#)

- ☐ **14. Aggregates in the temporal query language TQuel**
Snodgrass, R.T.; Gomez, S.; McKenzie, L.E., Jr.;
[Knowledge and Data Engineering, IEEE Transactions on](#)
Volume 5, Issue 5, Oct. 1993 Page(s):826 - 842
Digital Object Identifier 10.1109/69.243512
[AbstractPlus](#) | [Full Text: PDF\(1572 KB\)](#) IEEE JNL
[Rights and Permissions](#)

- ☐ **15. Automatic model-based semantic object extraction algorithm**
Jianping Fan; Xingquan Zhu; Lide Wu;
[Circuits and Systems for Video Technology, IEEE Transactions on](#)
Volume 11, Issue 10, Oct. 2001 Page(s):1073 - 1084
Digital Object Identifier 10.1109/76.954494
[AbstractPlus](#) | [References](#) | [Full Text: PDF\(360 KB\)](#) IEEE JNL
[Rights and Permissions](#)

- ☐ **16. Parallel Star Join+DataIndexes: efficient query processing in data warehouse**
Datta, A.; VanderMeer, D.; Ramamritham, K.;
[Knowledge and Data Engineering, IEEE Transactions on](#)
Volume 14, Issue 6, Nov.-Dec. 2002 Page(s):1299 - 1316
Digital Object Identifier 10.1109/TKDE.2002.1047769
[AbstractPlus](#) | [References](#) | [Full Text: PDF\(1973 KB\)](#) IEEE JNL
[Rights and Permissions](#)

- ☐ **17. A 2/sup d/-tree-based blocking method for microaggregating very large d**
Solanas, A.; Martinez-Balleste, A.; Domingo-Ferrer, J.; Mateo-Sanz, J.M.;
[Availability, Reliability and Security, 2006. ARES 2006. The First International](#)
20-22 April 2006 Page(s):7 pp.
Digital Object Identifier 10.1109/ARES.2006.1
[AbstractPlus](#) | [Full Text: PDF\(464 KB\)](#) IEEE CNF
[Rights and Permissions](#)

☐ Search Results

[BROWSE](#)

[SEARCH](#)

[IEEE XPLORE GUIDE](#)

Results for "(((aggregation<in>metadata) <and> (intervals<in>metadata))<and> (data..."

 e-mail

Your search matched 2 of 1408155 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

[View Session History](#)

[New Search](#)

Modify Search

(((aggregation<in>metadata) <and> (intervals<in>metadata))<and> (database<in>

[Search](#)

☐ Check to search only within this results set

Display Format: ☒ Citation ☐ Citation & Abstract

» Key

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

[view selected items](#)

[Select All](#) [Deselect All](#)

- ☐ 1. Large-sample and deterministic confidence intervals for online aggregation
Haas, P.J.;
[Scientific and Statistical Database Management, 1997. Proceedings., Ninth Int Conference on](#)
11-13 Aug. 1997 Page(s):51 - 62
Digital Object Identifier 10.1109/SSDM.1997.621151
[AbstractPlus](#) | Full Text: [PDF](#)(940 KB) IEEE CNF
[Rights and Permissions](#)
- ☐ 2. Computing temporal aggregates
Kline, N.; Snodgrass, R.T.;
[Data Engineering, 1995. Proceedings of the Eleventh International Conference](#)
6-10 March 1995 Page(s):222 - 231
Digital Object Identifier 10.1109/ICDE.1995.380389
[AbstractPlus](#) | Full Text: [PDF](#)(916 KB) IEEE CNF
[Rights and Permissions](#)



Terms used [sort aggregations](#)

Found 26,366 of 185,178

Sort results by

relevance



Save results to a Binder

Try an [Advanced Search](#)

Display results

expanded form



Search Tips

Try this search in [The ACM Guide](#)

☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Similarity queries I: Efficient similarity search and classification via rank aggregation](#)



Ronald Fagin, Ravi Kumar, D. Sivakumar

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: [pdf\(198.89 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We propose a novel approach to performing efficient similarity search and classification in high dimensional data. In this framework, the database elements are vectors in a Euclidean space. Given a query vector in the same space, the goal is to find elements of the database that are similar to the query. In our approach, a small number of independent "voters" rank the database elements based on similarity to the query. These rankings are then combined by a highly efficient aggregation algorithm. ...

2 [Exploiting early sorting and early partitioning for decision support query processing](#)

J. Claussen, A. Kemper, D. Kossmann, C. Wiesner

December 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 9 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: [pdf\(478.23 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

Decision support queries typically involve several joins, a grouping with aggregation, and/or sorting of the result tuples. We propose two new classes of query evaluation algorithms that can be used to speed up the execution of such queries. The algorithms are based on (1) *early sorting* and (2) *early partitioning*— or a combination of both. The idea is to push the sorting and/or the partitioning to the leaves, i.e., the base relations, of the query evaluation plans (QEPs) and ...

Keywords: Decision Support Systems, Early sorting and partitioning, Hash joins and hash teams, Performance evaluation, Query processing and optimization

3 [Incremental computation and maintenance of temporal aggregates](#)

Jun Yang, Jennifer Widom

October 2003 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 12 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(360.68 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citings](#), [index terms](#)

Abstract. We consider the problems of computing aggregation queries in temporal databases and of maintaining materialized temporal aggregate views efficiently. The latter problem is particularly challenging since a single data update can cause aggregate results to change over the entire time line. We introduce a new index structure called the *SB-tree*, which incorporates features from both *segment-trees* and *B-trees*. *SB-trees* support fast lookup of aggregate results based on ti ...

Keywords: Access methods, Aggregation, B-tree, Segment tree, Temporal database, View maintenance


4 On parallel processing of aggregate and scalar functions in object-relational DBMS



Michael Jaedicke, Bernhard Mitschang

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98**, Volume 27 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.43 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

Nowadays parallel object-relational DBMS are envisioned as the next great wave, but there is still a lack of efficient implementation concepts for some parts of the proposed functionality. Thus one of the current goals for parallel object-relational DBMS is to move towards higher performance. In this paper we develop a framework that allows to process user-defined functions with data parallelism. We will describe the class of partitionable functions that can be processed parallelly. We will ...

Keywords: aggregates, object-relational database systems, parallel query processing, user-defined functions

5 Query processing for relational data: Supporting ad-hoc ranking aggregates



Chengkai Li, Kevin Chen-Chuan Chang, Ihab F. Ilyas

June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press

Full text available:  [pdf\(344.23 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper presents a principled framework for efficient processing of ad-hoc *top-k* (ranking) aggregate queries, which provide the *k* groups with the highest aggregates as results. Essential support of such queries is lacking in current systems, which process the queries in a naïve materialize-group-sort scheme that can be prohibitively inefficient. Our framework is based on three fundamental principles. The Upper-Bound Principle dictates the requirements of early pruning, and ...

Keywords: OLAP, aggregate query, decision support, ranking, top-k query processing

6 Logics with aggregate operators



Lauri Hella, Leonid Libkin, Juha Nurmonen, Limsoon Wong

July 2001 **Journal of the ACM (JACM)**, Volume 48 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(323.27 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

We study adding aggregate operators, such as summing up elements of a column of a relation, to logics with counting mechanisms. The primary motivation comes from

database applications, where aggregate operators are present in all real life query languages. Unlike other features of query languages, aggregates are not adequately captured by the existing logical formalisms. Consequently, all previous approaches to analyzing the expressive power of aggregation were only capable of producing partial ...

Keywords: Aggregation, database, expressive power, locality, relational calculus

7 Query processing: Exploiting hierarchical clustering in evaluating multidimensional aggregation queries



Dimitri Theodoratos

November 2003 **Proceedings of the 6th ACM international workshop on Data warehousing and OLAP**

Publisher: ACM Press

Full text available: pdf(216.79 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Multidimensional aggregation queries constitute the single most important class of queries for data warehousing applications and decision support systems. The bottleneck in the evaluation of these queries is the join of the usually huge fact table with the restricted dimension tables (*star-join*). Recently, a multidimensional hierarchical clustering schema for star schemas is suggested. Subsequently, query evaluation plans for multidimensional queries appeared that essentially implement a ...

Keywords: multidimensional aggregation query, multidimensional hierarchical clustering, query transformations, star join

8 Aggregate predicate support in DBMS

Apostol (Paul) Natsev, Gene Y. C. Fuh, Weidong Chen, Chi-Huang Chiu, Jeffrey S. Vitter
January 2002 **Australian Computer Science Communications , Proceedings of the thirteenth Australasian database conference - Volume 5 ADC '02**, Volume 24 Issue 2

Publisher: Australian Computer Society, Inc., IEEE Computer Society Press

Full text available: pdf(1.57 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper we consider aggregate predicates and their support in database systems. Aggregate predicates are the predicate equivalent to aggregate functions in that they can be used to search for tuples that satisfy some aggregate property over a set of tuples (as opposed to simply computing an aggregate property over a set of tuples). The importance of aggregate predicates is exemplified by many modern applications that require ranked search, or top-*k* queries. Such queries are the norm ...

Keywords: aggregate predicates, nearest neighbor, query optimization

9 Online aggregation



Joseph M. Hellerstein, Peter J. Haas, Helen J. Wang

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: pdf(1.92 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Aggregation in traditional database systems is performed in batch mode: a query is submitted, the system processes a large volume of data over a long period of time, and, eventually, the final answer is returned. This archaic approach is frustrating to users and has been abandoned in most other areas of computing. In this paper we propose a new online aggregation interface that permits users to both observe the progress of their

aggregation queries and control execution on ...

10 Research session: query optimization #2: Optimizing nested queries with parameter sort orders

Ravindra Guravannavar, H. S. Ramanujam, S. Sudarshan

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

Full text available:  pdf(200.19 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Nested iteration is an important technique for query evaluation. It is the default way of executing nested subqueries in SQL. Although decorrelation often results in cheaper non-nested plans, decorrelation is not always applicable for nested subqueries. Nested iteration, if implemented properly, can also win over decorrelation for several classes of queries. Decorrelation is also hard to apply to nested iteration in user-defined SQL procedures and functions. Recent research has proposed evaluati ...

11 Visibility sorting and compositing without splitting for image layer decompositions



John Snyder, Jed Lengyel

July 1998 **Proceedings of the 25th annual conference on Computer graphics and interactive techniques**

Publisher: ACM Press

Full text available:  pdf(591.53 KB) Additional Information: [full citation](#), [references](#), [citing](#), [index terms](#)

Keywords: compositing, kd-tree, nonsplitting layered decomposition, occlusion cycle, occlusion graph, sprite, visibility sorting

12 Performance evaluation of the statistical aggregation by categorization in the SM3 system



C. K Baru, S. Y. W. Su

June 1984 **ACM SIGMOD Record , Proceedings of the 1984 ACM SIGMOD international conference on Management of data SIGMOD '84**, Volume 14 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.32 MB) Additional Information: [full citation](#), [abstract](#), [references](#)

To perform a statistical aggregation operation over a large file often requires that the records of the file be divided into categories based on the values of the attribute(s) over which some statistical computation is to be performed. It is rather inefficient to perform the necessary data transfer, categorization and statistical computation using a single processor. Parallel algorithms designed for multiprocessor systems have been proposed and their performance improvement over the conventional ...

13 Optimal aggregation algorithms for middleware



Ronald Fagin, Amnon Lotem, Moni Naor

May 2001 **Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available:  pdf(231.47 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

Assume that each object in a database has m grades, or scores, one for each of m attributes. For example, an object can have a color grade, that tells how red it is, and a shape grade, that tells how round it is. For each attribute, there is a sorted list, which lists each object and its grade under that attribute, sorted by grade (highest grade first).

There is some monotone *aggregation function*, or *combining rule*, such as min or average, that combines the individ ...

14 Research papers: stream aggregation: Semantics and evaluation techniques for window aggregates in data streams



Jin Li, David Maier, Kristin Tufte, Vassilis Papadimos, Peter A. Tucker

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(564.92 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

A windowed query operator breaks a data stream into possibly overlapping subsets of data and computes a result over each. Many stream systems can evaluate window aggregate queries. However, current stream systems suffer from a lack of an explicit definition of window semantics. As a result, their implementations unnecessarily confuse window definition with physical stream properties. This confusion complicates the stream system, and even worse, can hurt performance both in terms of memory usage ...

15 An NF2 relational interface for document retrieval, restructuring and aggregation



Kalervo Järvelin, Timo Niemi

July 1995 **Proceedings of the 18th annual international ACM SIGIR conference on Research and development in information retrieval**

Publisher: ACM Press

Full text available: pdf(985.40 KB) Additional Information: [full citation](#), [references](#), [citings](#), [index terms](#)

16 Fast algorithms for universal quantification in large databases



Goetz Graefe, Richard L. Cole

June 1995 **ACM Transactions on Database Systems (TODS)**, Volume 20 Issue 2

Publisher: ACM Press

Full text available: pdf(3.51 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#), [review](#)

Universal quantification is not supported directly in most database systems despite the fact that it adds significant power to a system's query processing and inference capabilities, in particular for the analysis of many-to-many relationships and of set-valued attributes. One of the main reasons for this omission has been that universal quantification algorithms and their performance have not been explored for large databases. In this article, we describe and compare three known algorithms ...

17 An array-based algorithm for simultaneous multidimensional aggregates



Yihong Zhao, Prasad M. Deshpande, Jeffrey F. Naughton

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: pdf(1.45 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

Computing multiple related group-bys and aggregates is one of the core operations of On-Line Analytical Processing (OLAP) applications. Recently, Gray et al. [GBLP95] proposed the "Cube" operator, which computes group-by aggregations over all possible subsets of the specified dimensions. The rapid acceptance of the importance of this operator has led to a variant of the Cube being proposed for the SQL standard. Several efficient algorithms for Relational OLAP (ROLAP) have been d ...

18 Parallel algorithms for the execution of relational database operations



Dina Bitton, Haran Boral, David J. DeWitt, W. Kevin Wilkinson
September 1983 **ACM Transactions on Database Systems (TODS)**, Volume 8 Issue 3

Publisher: ACM Press

Full text available: pdf(2.07 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents and analyzes algorithms for parallel processing of relational database operations in a general multiprocessor framework. To analyze alternative algorithms, we introduce an analysis methodology which incorporates I/O, CPU, and message costs and which can be adjusted to fit different multiprocessor architectures. Algorithms are presented and analyzed for sorting, projection, and join operations. While some of these algorithms have been presented and analyzed previously, we ...

Keywords: aggregate operations, database machines, join operation, parallel processing, projection operator, sorting

19 Query evaluation techniques for large databases



Goetz Graefe
June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(9.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

20 Object-based and image-based object representations



Hanan Samet
June 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 2

Publisher: ACM Press

Full text available: pdf(1.05 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An overview is presented of object-based and image-based representations of objects by their interiors. The representations are distinguished by the manner in which they can be used to answer two fundamental queries in database applications: (1) Feature query: given an object, determine its constituent cells (i.e., their locations in space). (2) Location query: given a cell (i.e., a location in space), determine the identity of the object (or objects) of which it is a member as well as the re ...

Keywords: Access methods, R-trees, feature query, geographic information systems (GIS), image space, location query, object space, octrees, pyramids, quadrees, space-filling curves, spatial databases

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)Search: ☒ The ACM Digital Library ☐ The Guide**SEARCH**

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)Terms used **multiple database records aggregations**Found **52,501** of **185,178**Sort results
by ☒Display
results ☒ [Save results to a Binder](#) [Search Tips](#)☐ Open results in a new
windowTry an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐**1** [Session: database and program conversion: Towards the support of integrated views](#)[of multiple databases: an aggregate schema facility](#)

Donald Swartwout, James P. Fry

May 1978 **Proceedings of the 1978 ACM SIGMOD international conference on
management of data****Publisher:** ACM PressFull text available: [pdf\(1.03 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#)

Supporting multiple user views of databases is currently an important problem area in database management system development. An interesting facet of this problem arises whenever a user needs an integrated view of several distinct databases. Using traditional database concepts, an aggregate schema facility has been developed to address this problem. The basic functions of an aggregate schema facility are discussed, as well as their implementation in a CODASYL/DBTG-like environment. Interest in a ...

Keywords: aggregate schema, data definition languages, data translation, database integration, database management systems, database restructuring, distributed databases, dynamic translation

2 [Automatic high-quality reengineering of database programs by abstraction,](#)[transformation and reimplementation](#)

Yossi Cohen, Yishai A. Feldman

July 2003 **ACM Transactions on Software Engineering and Methodology (TOSEM),
Volume 12 Issue 3****Publisher:** ACM PressFull text available: [pdf\(245.97 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Old-generation database models, such as the indexed-sequential, hierarchical, or network models, provide record-level access to their data, with all application logic residing in the hosting program. In contrast, relational databases can perform complex operations, such as filter, aggregation, and join, on multiple records without an external specification of the record-access logic. Programs written for relational databases attempt to move as much of the application logic as possible into the d ...

Keywords: Database program reengineering, query graphs, temporal abstraction, the plan calculus

3 The Logical Record Access Approach to Database Design



Toby J. Teorey, James P. Fry

June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2

Publisher: ACM Press

Full text available: pdf(2.81 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

4 An overview of data warehousing and OLAP technology



Surajit Chaudhuri, Umeshwar Dayal

March 1997 **ACM SIGMOD Record**, Volume 26 Issue 1

Publisher: ACM Press

Full text available: pdf(101.60 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Data warehousing and on-line analytical processing (OLAP) are essential elements of decision support, which has increasingly become a focus of the database industry. Many commercial products and services are now available, and all of the principal database management system vendors now have offerings in these areas. Decision support places some rather different requirements on database technology compared to traditional on-line transaction processing applications. This paper provides an overview ...

5 Techniques for Structuring Database Records



Salvatore T. March

March 1983 **ACM Computing Surveys (CSUR)**, Volume 15 Issue 1

Publisher: ACM Press

Full text available: pdf(3.02 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

6 Query evaluation techniques for large databases



Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(9.37 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

7 Progressive evaluation of nested aggregate queries

Kian-Lee Tan, Cheng Hian Goh, Beng Chin Ooi

December 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 9 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(380.81 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

In many decision-making scenarios, decision makers require rapid feedback to their queries, which typically involve aggregates. The traditional *blocking execution model* can no longer meet the demands of these users. One promising approach in the literature, called *online aggregation*, evaluates an aggregation query progressively as follows: as soon as certain data have been evaluated, approximate answers are produced with their respective running confidence intervals; as more data a ...

Keywords: Approximate answers, Multi-threading, Nested aggregate queries, Online aggregation, Progressive query processing

8 Aggregation everywhere: data reduction and transformation in the Phoenix data warehouse



Steven Tolkin

November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP**

Publisher: ACM Press

Full text available: pdf(1.23 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the Phoenix system, which loads a data warehouse and then reports against it. Between the raw atomic data of the source system and the business measures presented to users there are many computing environments. Aggregation occurs everywhere: initial bucketing by the natural keys on the mainframe, loading the fact table using a mapping table, maintaining aggregate tables and reporting tables in the data base, in the GUI, in SQL queries issued on behalf of client tools by ...

Keywords: OLAP, SQL, aggregation, data lineage, data warehouse

9 Similarity queries I: Efficient similarity search and classification via rank aggregation



Ronald Fagin, Ravi Kumar, D. Sivakumar

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(198.89 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We propose a novel approach to performing efficient similarity search and classification in high dimensional data. In this framework, the database elements are vectors in a Euclidean space. Given a query vector in the same space, the goal is to find elements of the database that are similar to the query. In our approach, a small number of independent "voters" rank the database elements based on similarity to the query. These rankings are then combined by a highly efficient aggregation algorithm. ...

10 Data warehousing: Integrating compression and execution in column-oriented database systems



Daniel Abadi, Samuel Madden, Miguel Ferreira

June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press

Full text available: pdf(291.42 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Column-oriented database system architectures invite a re-evaluation of how and when data in databases is compressed. Storing data in a column-oriented fashion greatly increases the similarity of adjacent records on disk and thus opportunities for compression. The ability to compress many adjacent tuples at once lowers the per-tuple cost of compression, both in terms of CPU and space overheads. In this paper, we discuss

how we extended C-Store (a column-oriented DBMS) with a compression sub-system ...

Keywords: column-oriented databases, column-stores, database compression, query execution

11 A Gopher interface to relational databases



Paul Lindner

November 1993 **Proceedings of the 21st annual ACM SIGUCCS conference on User services**

Publisher: ACM Press

Full text available: pdf(446.64 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

12 Tools & techniques track: frameworks for building libraries: Using collection descriptions to enhance an aggregation of harvested item-level metadata



Muriel Foulonneau, Timothy W. Cole, Thomas G. Habing, Sarah L. Shreeves

June 2005 **Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries**

Publisher: ACM Press

Full text available: pdf(1.09 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As an increasing number of digital library projects embrace the harvesting of item-level descriptive metadata, issues of description granularity and concerns about potential loss of context when harvesting item-level metadata take on greater significance. Collection-level description can provide valuable context for item-level metadata records harvested from disparate and heterogeneous providers. This paper describes an ongoing experiment using collection-level description in concert with item-I ...

Keywords: collection-level description, descriptive metadata, metadata aggregation, open archives initiative

13 Data streams: On-the-fly sharing for streamed aggregation



Sailesh Krishnamurthy, Chung Wu, Michael Franklin

June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press

Full text available: pdf(1.11 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Data streaming systems are becoming essential for monitoring applications such as financial analysis and network intrusion detection. These systems often have to process many similar but different queries over common data. Since executing each query separately can lead to significant scalability and performance problems, it is vital to share resources by exploiting similarities in the queries. In this paper we present ways to efficiently share streaming aggregate queries with differing periodic ...

Keywords: aggregation, multiple-query optimization, shared processing, streaming data

14 The HiPAC project: combining active databases and timing constraints



M. J. Carey, M. Livny, R. Jauhari

March 1988 **ACM SIGMOD Record**, Volume 17 Issue 1

Publisher: ACM Press

Full text available: pdf(1.39 MB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The HiPAC (High Performance ACTIVE database system) project addresses two critical

problems in time-constrained data management: the handling of timing constraints in databases, and the avoidance of wasteful polling through the use of situation-action rules that are an integral part of the database and are monitored by DBMS's condition monitor. A rich knowledge model provides the necessary primitives for definition of timing constraints, situation-action rules, and precipitating events. The ...

15 A J2EE application for process accounting, LPAR accounting, and transaction accounting



C. Eric Wu, William P. Horn

July 2005 **Proceedings of the 5th international workshop on Software and performance WOSP '05**

Publisher: ACM Press

Full text available: pdf(751.27 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Accounting is critical for information technology budgeting and chargeback. Traditional accounting in UNIX/Linux systems is known as process accounting, in which an accounting record is created when a process ends. System administrators typically aggregate accounting records based on individual users or groups. As Web and application servers along with databases handle requests and transactions for multiple entities in various Web applications and services, LPAR accounting and transaction account ...

Keywords: ARM transactions, process accounting, project accounting, resource usage, transaction accounting

16 Authentication and integrity in outsourced databases



Einar Mykletun, Maithili Narasimha, Gene Tsudik

May 2006 **ACM Transactions on Storage (TOS)**, Volume 2 Issue 2

Publisher: ACM Press

Full text available: pdf(531.47 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In the Outsourced Database (ODB) model, entities outsource their data management needs to a third-party service provider. Such a service provider offers mechanisms for its clients to create, store, update, and access (query) their databases. This work provides mechanisms to ensure data integrity and authenticity for outsourced databases. Specifically, this article provides mechanisms that assure the querier that the query results have not been tampered with and are authentic (with respect to the ...

Keywords: Outsourced databases, authentication, data authenticity, data integrity, integrity, signature aggregation, storage

17 Research sessions: non-standard query processing: Fast computation of database operations using graphics processors



Naga K. Govindaraju, Brandon Lloyd, Wei Wang, Ming Lin, Dinesh Manocha

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(386.13 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

We present new algorithms for performing fast computation of several common database operations on commodity graphics processors. Specifically, we consider operations such as conjunctive selections, aggregations, and semi-linear queries, which are essential computational components of typical database, data warehousing, and data mining applications. While graphics processing units (GPUs) have been designed for fast display of geometric primitives, we utilize the inherent pipelining and parallelism ...

Keywords: aggregation, graphics processor, query optimization, selection query, selectivity analysis, semi-linear query

18 GPGPU: general purpose computation on graphics hardware



David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 **Proceedings of the conference on SIGGRAPH 2004 course notes SIGGRAPH '04**

Publisher: ACM Press

Full text available: pdf(63.03 MB) Additional Information: [full citation](#), [abstract](#)

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

19 Research sessions: indexing and tuning: Transaction support for indexed summary views



Goetz Graefe, Michael Zwilling

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(168.70 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

Materialized views have become a standard technique for performance improvement in decision support databases and for a variety of monitoring purposes. In order to avoid inconsistencies and thus unpredictable query results, materialized views and their indexes should be maintained immediately within user transaction just like indexes on ordinary tables. Unfortunately, the smaller a materialized view is, the higher the concurrency contention between queries and updates as well as among concurrent ...

20 Tools and transformations—rigorous and otherwise—for practical database design



Arnon Rosenthal, David Reiner

June 1994 **ACM Transactions on Database Systems (TODS)**, Volume 19 Issue 2

Publisher: ACM Press

Full text available: pdf(3.19 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We describe the tools and theory of a comprehensive system for database design, and show how they work together to support multiple conceptual and logical design processes. The Database Design and Evaluation Workbench (DDEW) system uses a rigorous, information-content-preserving approach to schema transformation, but combines it with heuristics, guess work, and user interactions. The main contribution lies in illustrating how theory was adapted to a practical system, and how the consistency ...

Keywords: applications of database theory, computer-aided software engineering, data model translation, database design, database equivalence, design heuristics, entity-relationship model, heuristics, normalization, view integration

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

parittions records aggregations

SEARCH

THE ACM DIGITAL LIBRARY



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [parittions](#) [records](#) [aggregations](#)

Found 7,778 of 185,178

Sort results
by

relevance



Save results to a Binder

Display
results

expanded form



Search Tips

☐ Open results in a new
window

Try an [Advanced Search](#)

Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [A robust, optimization-based approach for approximate answering of aggregate queries](#)



Surajit Chaudhuri, Gautam Das, Vivek Narasayya

May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01**, Volume 30 Issue 2

Publisher: ACM Press

Full text available: pdf(221.91 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The ability to approximately answer aggregation queries accurately and efficiently is of great benefit for decision support and data mining tools. In contrast to previous sampling-based studies, we treat the problem as an *optimization* problem whose goal is to minimize the error in answering queries in the given workload. A key novelty of our approach is that we can tailor the choice of samples to be robust even for workloads that are "similar" but not necessarily identical ...

2 [Session: database and program conversion: Towards the support of integrated views of multiple databases: an aggregate schema facility](#)



Donald Swartwout, James P. Fry

May 1978 **Proceedings of the 1978 ACM SIGMOD international conference on management of data**

Publisher: ACM Press

Full text available: pdf(1.03 MB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Supporting multiple user views of databases is currently an important problem area in database management system development. An interesting facet of this problem arises whenever a user needs an integrated view of several distinct databases. Using traditional database concepts, an aggregate schema facility has been developed to address this problem. The basic functions of an aggregate schema facility are discussed, as well as their implementation in a CODASYL/DBTG-like environment. Interest in a ...

Keywords: aggregate schema, data definition languages, data translation, database integration, database management systems, database restructuring, distributed databases, dynamic translation

3 [Optimizing spatial Min/Max aggregations](#)

Donghui Zhang, J. Tsotras

Aggregate computation over a collection of spatial objects appears in many real-life applications. Aggregates are computed on values (weights) associated with spatial objects, for example, the temperature or rainfall over the area covered by the object. In this paper we concentrate on MIN/MAX aggregations: "given a query rectangle, find the minimum/maximum weight among all objects intersecting the query rectangle." Traditionally such queries have been performed as range searches. A ...

Keywords: Indexing, Min/Max, Spatial aggregates


4 Spatial Query Processing Algorithms: Improving min/max aggregation over spatial objects



Donghui Zhang, Vassilis J. Tsotras

November 2001 **Proceedings of the 9th ACM international symposium on Advances in geographic information systems**

Publisher: ACM Press

Full text available:  [pdf\(1.70 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We examine the problem of computing MIN/MAX aggregates over a collection of spatial objects. Each spatial object is associated with a weight (value), for example, the average temperature or rainfall over the area covered by the object. Given a query rectangle, the MIN/MAX problem computes the minimum/maximum weight among all objects intersecting the query rectangle. Traditionally such queries have been performed as range search queries. Assuming that the objects are indexed by a spatial access m ...

Keywords: Min/Max, indexing, spatial aggregates

5 Physical interface: TAG: a Tiny AGgregation service for ad-hoc sensor networks



Samuel Madden, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Publisher: ACM Press

Full text available:  [pdf\(2.19 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

We present the Tiny AGgregation (TAG) service for aggregation in low-power, distributed, wireless environments. TAG allows users to express simple, declarative queries and have them distributed and executed efficiently in networks of low-power, wireless sensors. We discuss various generic properties of aggregates, and show how those properties affect the performance of our in network approach. We include a performance study demonstrating the advantages of our approach over traditional centralize ...

6 Techniques for Structuring Database Records



Salvatore T. March

March 1983 **ACM Computing Surveys (CSUR)**, Volume 15 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(3.02 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

7 Complete answer aggregates for treelike databases: a novel approach to combine querying and navigation





Holger Meuss, Klaus U. Schulz

April 2001 **ACM Transactions on Information Systems (TOIS)**, Volume 19 Issue 2

Publisher: ACM Press

Full text available: pdf(356.60 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The use of markup languages like SGML, HTML or XML for encoding the structure of documents or linguistic data has led to many databases where entries are adequately described as trees. In this context querying formalisms are interesting that offer the possibility to refer both to textual content and logical structure. We consider models where the structure specified in a query is not only used as a filter, but also for selecting and presenting different parts of the data. If answers are formalized ...

Keywords: SGML, XML, answer presentation, information retrieval, logic, query languages, semistructured data, structured documents, tree databases, tree matching

8 Tools & techniques track: frameworks for building libraries: Using collection descriptions to enhance an aggregation of harvested item-level metadata



Muriel Foulonneau, Timothy W. Cole, Thomas G. Habing, Sarah L. Shreeves

June 2005 **Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries**

Publisher: ACM Press

Full text available: pdf(1.09 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As an increasing number of digital library projects embrace the harvesting of item-level descriptive metadata, issues of description granularity and concerns about potential loss of context when harvesting item-level metadata take on greater significance. Collection-level description can provide valuable context for item-level metadata records harvested from disparate and heterogeneous providers. This paper describes an ongoing experiment using collection-level description in concert with item-level ...

Keywords: collection-level description, descriptive metadata, metadata aggregation, open archives initiative

9 Usage and relationships: An architecture for the aggregation and analysis of scholarly usage data



Johan Bollen, Herbert Van de Sompel

June 2006 **Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries JCDL '06**

Publisher: ACM Press

Full text available: pdf(753.90 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Although recording of usage data is common in scholarly information services, its exploitation for the creation of value-added services remains limited due to concerns regarding, among others, user privacy, data validity, and the lack of accepted standards for the representation, sharing and aggregation of usage data. This paper presents a technical, standards-based architecture for sharing usage information, which we have designed and implemented. In this architecture, OpenURL-compliant linking ...

Keywords: OAI-PMH, aggregation, analysis, architecture, digital libraries, openURL, standards, usage data

10 Aggregate structure identification and its application to program analysis



G. Ramalingam, John Field, Frank Tip

January 1999 **Proceedings of the 26th ACM SIGPLAN-SIGACT symposium on**



Principles of programming languages

Publisher: ACM Press

Full text available: [pdf\(1.92 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

11 Performance evaluation of the statistical aggregation by categorization in the SM3 system



C. K Baru, S. Y. W. Su

June 1984 **ACM SIGMOD Record , Proceedings of the 1984 ACM SIGMOD international conference on Management of data SIGMOD '84**, Volume 14 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.32 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

To perform a statistical aggregation operation over a large file often requires that the records of the file be divided into categories based on the values of the attribute(s) over which some statistical computation is to be performed. It is rather inefficient to perform the necessary data transfer, categorization and statistical computation using a single processor. Parallel algorithms designed for multiprocessor systems have been proposed and their performance improvement over the conventional ...

12 Research papers: stream aggregation: Multiple aggregations over data streams



Rui Zhang, Nick Koudas, Beng Chin Ooi, Divesh Srivastava

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: [pdf\(403.02 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

Monitoring aggregates on IP traffic data streams is a compelling application for data stream management systems. The need for exploratory IP traffic data analysis naturally leads to posing related aggregation queries on data streams, that differ only in the choice of grouping attributes. In this paper, we address this problem of efficiently computing multiple aggregations over high speed data streams, based on a two-level LFTA/HFTA DSMS architecture, inspired by Gigascope. Our first contribution ...

13 Online aggregation



Joseph M. Hellerstein, Peter J. Haas, Helen J. Wang

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.92 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Aggregation in traditional database systems is performed in batch mode: a query is submitted, the system processes a large volume of data over a long period of time, and, eventually, the final answer is returned. This archaic approach is frustrating to users and has been abandoned in most other areas of computing. In this paper we propose a new online aggregation interface that permits users to both observe the progress of their aggregation queries and control execution on ...

14 Technical poster session 1: multimedia analysis, processing, and retrieval:



Calculation of an aggregated level of interest function for recorded events

Rahul Nair

October 2004 **Proceedings of the 12th annual ACM international conference on Multimedia**

Publisher: ACM Press

Full text available:

Additional Information:

As recording technology becomes pervasive there is a dramatic increase in the number of events being recorded in multimedia. The challenge now facing users is to quickly view the recorded content in the least amount of time. While there are several methods to analyze video based on ambient noise, scene changes, slide transitions, etc..., these techniques merely find features in the recording, they do not reveal which sections are important.

This paper presents a method to calculate a L ...

Keywords: bookmark aggregation, level of interest, multimedia, skimming, video browsing, visualization


15 Historical spatio-temporal aggregation



Yufei Tao, Dimitris Papadias

January 2005 **ACM Transactions on Information Systems (TOIS)**, Volume 23 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.42 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Spatio-temporal databases store information about the positions of individual objects over time. However, in many applications such as traffic supervision or mobile communication systems, only summarized data, like the number of cars in an area for a specific period, or phone-calls serviced by a cell each day, is required. Although this information can be obtained from operational databases, its computation is expensive, rendering online processing inapplicable. In this paper, we present special ...

Keywords: Aggregation, access methods, cost models


16 Compressed data cubes for OLAP aggregate query approximation on continuous dimensions



Jayavel Shanmugasundaram, Usama Fayyad, P. S. Bradley

August 1999 **Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining**

Publisher: ACM Press

Full text available:  pdf(1.12 MB)

Additional Information: [full citation](#), [references](#), [citings](#), [index terms](#)


Keywords: OLAP, approximate query answering, clustering, data cubes, data mining, density estimation

17 Progressive evaluation of nested aggregate queries

Kian-Lee Tan, Cheng Hian Goh, Beng Chin Ooi

December 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 9 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(380.81 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

In many decision-making scenarios, decision makers require rapid feedback to their queries, which typically involve aggregates. The traditional *blocking execution model* can no longer meet the demands of these users. One promising approach in the literature, called *online aggregation*, evaluates an aggregation query progressively as follows: as soon as certain data have been evaluated, approximate answers are produced with their

respective running confidence intervals; as more data a ...

Keywords: Approximate answers, Multi-threading, Nested aggregate queries, Online aggregation, Progressive query processing

18 Supporting education: Metadata aggregation and "automated digital libraries": a retrospective on the NSDL experience



Carl Lagoze, Dean Krafft, Tim Cornwell, Naomi Dushay, Dean Eckstrom, John Saylor
June 2006 **Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries JCDL '06**

Publisher: ACM Press

Full text available: pdf(346.87 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Over three years ago, the Core Integration team of the National Science Digital Library (NSDL) implemented a digital library based on metadata aggregation using Dublin Core and OAI-PMH. The initial expectation was that such low-barrier technologies would be relatively easy to automate and administer. While this architectural choice permitted rapid deployment of a production NSDL, our three years of experience have contradicted our original expectations of easy automation and low people cost. We ...

Keywords: NSDL, OAI-PMH, architecture, interoperability, metadata

19 Processing time-constrained aggregate queries in CASE-DB



Wen-Chi Hou, Gultekin Ozsoyoglu
June 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 2

Publisher: ACM Press

Full text available: pdf(2.62 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In this paper, we present an algorithm to strictly control the time to process an estimator for an aggregate relational query. The algorithm implemented in a prototype database management system, called CASE-DB, iteratively samples from input relations, and evaluates the associated estimator until the time quota expires. In order to estimate the time cost of a query, CASE-DB uses adaptive time cost formulas. The formulas are adaptive in that the parameters of the formulas can be ...

Keywords: estimation, relational algebra, risk of overspending, sampling, selectivity, time constraints

20 The Logical Record Access Approach to Database Design



Toby J. Teorey, James P. Fry
June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2

Publisher: ACM Press

Full text available: pdf(2.81 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)

RESULT LIST

15 results found in the Worldwide database for:

relational in the title AND **aggregation** in the title or abstract

(Results are sorted by date of upload in database)

1 Partial pre-aggregation in relational database queries

Inventor: LARSON PER-AKE (US); GALINDO-LEGARIA CESAR A (US) Applicant: MICROSOFT CORP (US)

EC:

IPC: **G06F7/00; G06F7/00**; (IPC1-7): G06F7/00

Publication info: **US2005240556** - 2005-10-27

2 REAL-TIME AGGREGATION OF UNSTRUCTURED DATA INTO STRUCTURED DATA FOR SQL PROCESSING BY A RELATIONAL DATABASE ENGINE

Inventor: CHOI ARTHUR; LEYBA TODD; (+2)

Applicant: IBM (US); IBM FRANCE (FR)

EC: G06F17/30H

IPC: **G06F17/30; G06F17/30**; (IPC1-7): G06F17/30

Publication info: **WO2004084092** - 2004-09-30

3 REMOTE SCORING AND AGGREGATING SIMILARITY SEARCH ENGINE FOR USE WITH RELATIONAL DATABASES

Inventor: RIPLEY JOHN R; ANANTHA RAM; (+1)

Applicant: INFOGLIDE SOFTWARE CORP (US)

EC:

IPC: **G06F17/30; G06F17/30**; (IPC1-7): G06F17/30

Publication info: **WO2005017778** - 2005-02-24

4 Remote scoring and aggregating similarity search engine for use with relational databases

Inventor: RIPLEY JOHN R (US); ANANTHA RAM (US); (+1) Applicant:

EC: G06F17/30H

IPC: **G06F17/30; G06F17/30**; (IPC1-7): G06F7/00

Publication info: **US2004064449** - 2004-04-01

5 Querying encrypted data in a relational database system

Inventor: HACIGUMUS VAHIT HAKAN (US); IYER BALAKRISHNA RAGHAVENDRA (US); (+1)

Applicant: IBM (US)

EC: G06F17/30H

IPC: **G06F12/14; H04L9/00; G06F12/14** (+2)

Publication info: **US2004243816** - 2004-12-02

6 Multi-part looked-up table fields and its use in data processing operations involving multiple tables of a relational database

Inventor: BELLEW MATTHEW A (US)

Applicant:

EC: G06F17/30H

IPC: **G06F12/00; G06F17/30; G06F12/00** (+2)

Publication info: **US2003131215** - 2003-07-10

7 Aggregate join index for relational databases

Inventor: AU GRACE KWAN-ON (US); HOANG CHI KIM (US) Applicant: NCR INT INC (US)

EC: G06F17/30S1R

IPC: **G06F17/30; G06F17/30**; (IPC1-7): G06F17/30

Publication info: **EP1164509** - 2001-12-19

8 Active caching for multi-dimensional data sets in relational database management system

Inventor: DESHPANDE PRASAD MANIKARAO (US); RAMASAMY KARTHIKEYAN (US); (+2)

Applicant: NCR CORP (US)

EC: G06F17/30H

IPC: **G06F17/30; G06F17/30**; (IPC1-7): G06F17/30

Publication info: **US6601062** - 2003-07-29

9 Method for determining the computability of data for an active multi-dimensional cache in a relational database management system

Inventor: DESHPANDE PRASAD MANIKARAO (US); RAMASAMY KARTHIKEYAN (US); (+2)

Applicant: NCR CORP (US)

EC: G06F17/30H6; G06F17/30S1

IPC: **G06F17/30; G06F17/30**; (IPC1-7): G06F17/30

Publication info: **US6763357** - 2004-07-13

10 Network management event storage and manipulation using relational database technology

Inventor: COSTA BRIAN (US); DAS CHRIS P (US);
(+1)

Applicant: HEWLETT PACKARD CO (US)

EC: H04L12/24A6; H04L12/24D

IPC: **G06F15/00; G06F17/30; H04L12/24** (+5)

Publication info: **EP0961439** - 1999-12-01

Data supplied from the *esp@cenet* database - Worldwide

RESULT LIST

15 results found in the Worldwide database for:
transactions in the title AND **aggregation** in the title or abstract
(Results are sorted by date of upload in database)

11 System and method for syndicated transactions

Inventor: FREISHTAT GREGG (US); RIJSINGHANI
VIKAS (US)

Applicant:

EC: G06Q30/00C

IPC: **G06Q30/00**; **G06Q30/00**; (IPC1-7): G06F17/60
(+3)

Publication info: **US2001037294** - 2001-11-01

12 Apparatus and method for automated aggregation and delivery of and transactions involving electronic personal information or data apparatus

Inventor: FREISHTAT GREGG; RAJAN PALANISWAMY

Applicant: VERTICALONE CORP

EC:

IPC: **G06F15/16**; **G06F17/30**; **G06F15/16** (+4)

Publication info: **AU4721000** - 2000-09-28

13 Apparatus and methods for automated aggregation and delivery of and transactions involving electronic personal information or data

Inventor: FREISHTAT GREGG (US); RAJAN
PALANISWAMY (US)

Applicant: VERTICALONE CORP (US)

EC: G06Q30/00; H04L29/06; (+8)

IPC: **G06Q30/00**; **H04L29/06**; **H04L29/08** (+4)

Publication info: **US6317783** - 2001-11-13

14 Apparatus and method for automated aggregation and delivery of and transactions involving electronic personal information or data

Inventor: FREISHTAT GREGG (US); RAJAN
PALANISWAMY (US)

Applicant: VERTICALONE CORP (US)

EC: G06F21/00N9A2P; G06F21/00N9A2P2; (+1)

IPC: **G06F17/30**; **G06F21/00**; **G06F17/30** (+2)

Publication info: **EP1107125** - 2001-06-13

15 APPARATUS AND METHOD FOR AUTOMATED AGGREGATION AND DELIVERY OF AND TRANSACTIONS INVOLVING ELECTRONIC PERSONAL INFORMATION OR DATA

Inventor: BURSON ROBERT (US); RAJAN
PALANISWAMY (US); (+4)

Applicant: VERTICALONE CORP (US)

EC:

IPC: **G06F17/40**; **H04L12/16**; **G06F17/40** (+3)

Publication info: **CA2308242** - 2000-05-04

Data supplied from the **esp@cenet** database - Worldwide

RESULT LIST

5 results found in the Worldwide database for:
databases in the title AND **partition** in the title or abstract
(Results are sorted by date of upload in database)

1 SYSTEM AND METHOD FOR RECOVERY UNITS IN DATABASES

Inventor: BYRNE PETER (US)

Applicant: MICROSOFT CORP (US); BYRNE PETER (US)

EC:

IPC: G06F17/30; G06F17/30

Publication info: WO2005091736 - 2005-10-06

2 Partition boundary determination using random sampling on very large databases

Inventor: HARPER JOHN WILLIAM (US); SLISHMAN
GORDON ROBERT (US)

Applicant: IBM (US)

EC: G06F17/30S1

IPC: (IPC1-7): G06F7/00

Publication info: US2003004944 - 2003-01-02

3 A method to efficiently partition large hyperlinked databases by hyperlink structure

Inventor: FLAKE GARY WILLIAM (US)

Applicant: NIPPON ELECTRIC CO (JP)

EC: G06F17/30G4

IPC: G06F17/30; G06F17/30; (IPC1-7): G06F17/30

Publication info: EP1076299 - 2001-02-14

4 On-line reorganization in object-oriented databases

Inventor: LAKHAMRAJU MOHANA KRISHNA (US);
RASTOGI RAJEEV (US); (+2)

Applicant: LUCENT TECHNOLOGIES INC (US)

EC: G06F17/30B; G06F17/30S3

IPC: G06F17/30; G06F17/30; (IPC1-7): G06F17/30
(+1)

Publication info: US6343296 - 2002-01-29

5 Query processor for parallel processing in homogenous and heterogenous databases

Inventor: COHEN GERALD D (US); STOUT RALPH L
(US); (+1)

Applicant: INFORMATION BUILDERS INC (US)

EC: G06F17/30N; G06F17/30T

IPC: G06F17/30; G06F17/30; (IPC1-7): G06F17/30
(+1)

Publication info: US5590319 - 1996-12-31

Data supplied from the *esp@cenet* database - Worldwide

RESULT LIST

15 results found in the Worldwide database for:
transactions in the title AND **aggregation** in the title or abstract
(Results are sorted by date of upload in database)

- 1 Detecting structuring of financial transactions**
Inventor: STRETTON PETER J (GB) Applicant: IBM (US)
EC: IPC: **G06Q40/00; G06F11/14; G06F11/14** (+1)
Publication info: **US2006095368** - 2006-05-04
- 2 Method for a variable rebate tier structure for card transactions**
Inventor: MITCHELL ERICA L (US); FREUD ALIZA (US); Applicant:
(+2)
EC: G06Q30/00A IPC: **G06Q30/00; G06Q30/00; (IPC1-7): G06F17/60**
Publication info: **US2005065850** - 2005-03-24
- 3 Method for a variable rebate tier structure for card transactions**
Inventor: MITCHELL ERICA L (US); FREUD ALIZA (US); Applicant:
(+2)
EC: G06Q30/00A IPC: **G06Q30/00; G06Q30/00; (IPC1-7): G06F17/60**
Publication info: **US2005065849** - 2005-03-24
- 4 Method for a variable rebate tier structure for card transactions**
Inventor: MITCHELL ERICA L (US); FREUD ALIZA (US); Applicant:
(+2)
EC: G06Q30/00A IPC: **G06Q30/00; G06Q30/00; (IPC1-7): G06F17/60**
Publication info: **US2005065848** - 2005-03-24
- 5 Method and system to facilitate payments to satisfy payment obligations resulting from purchase transactions**
Inventor: CHEN ANDREW D (US); PHILLIPS BRIAN A Applicant:
(US); (+4)
EC: G06Q20/00K3B; G06Q30/00B IPC: **G06Q20/00; G06Q30/00; G06Q20/00** (+3)
Publication info: **US2005097040** - 2005-05-05
- 6 Rule based aggregation of files and transactions in a switched file system**
Inventor: MILOUSHEV VLADIMIR (US); NICKOLOV Applicant: FORCE COMMUNICATIONS INC Z (US)
PETER (US)
EC: G06F17/30F; G06F17/30N; (+3) IPC: **G06F17/30; H04L29/06; H04L29/08** (+5)
Publication info: **US2004133577** - 2004-07-08
- 7 Method for a variable rebate tier structure for card transactions**
Inventor: MITCHELL ERICA L (US); FREUD ALIZA (US); Applicant:
(+2)
EC: G06Q30/00A IPC: **G06Q30/00; G06Q30/00; (IPC1-7): G06F17/60**
Publication info: **US2004068438** - 2004-04-08
- 8 SYSTEM, METHOD AND MEDIUM FOR FACILITATING TRANSACTIONS OVER A NETWORK**
Inventor: FROST COLIN (GB) Applicant: 2020ME HOLDINGS LTD (GB); FROST COLIN
(GB)
EC: G06Q30/00C; H04L29/06; (+2) IPC: **G06Q30/00; H04L29/06; H04L29/08** (+4)
Publication info: **WO0205153** - 2002-01-17
- 9 Sales transactions for transfer of agricultural products**
Inventor: DINES DAVID (US); TRACY MARK (US); (+3) Applicant:
EC: G06Q30/00C IPC: **G06Q30/00; G06Q30/00; (IPC1-7): G06F17/60**
Publication info: **US2002052795** - 2002-05-02
- 10 Sales transactions for transfer of commodities**
Inventor: DINES DAVID (US); TRACY MARK (US); (+3) Applicant:

EC: G06Q30/00C

IPC: G06Q30/00; G06Q30/00; (IPC1-7): G06F17/60

Publication info: **US2002052793** - 2002-05-02

Data supplied from the *esp@cenet* database - Worldwide

RESULT LIST

2 results found in the Worldwide database for:
aggregates in the title AND **partitions** in the title or abstract
(Results are sorted by date of upload in database)

1 Container for storing building aggregates

Inventor: LAMOND KEITH ARTHUR

Applicant: LAMOND KEITH ARTHUR

EC: B28C7/00B1C

IPC: **B28C7/00; B28C7/00**; (IPC1-7): B28C9/02
(+1)

Publication info: **GB2211829** - 1989-07-12

2 ZPUSOB VYROBY PREPAZEK TRUBKOVYCH SVAZKU TEPLOVYMENNYCH AGREGATU ZE SYNTETICKE PRYSKYRICE

Inventor: JERABEK JAROMIR (SA)

Applicant: JERABEK JAROMIR (SA)

EC:

IPC: **B29C39/02; B29C39/02**; (IPC1-7): B29C39/02

Publication info: **CS8506074** - 1988-01-15

Data supplied from the **esp@cenet** database - Worldwide